Numerical Weather Prediction

Digital Filtering and IAU

Forecasting Research Technical Report No. 427

Adam M. Clayton

email: nwp_publications@metoffice.com

©Crown Copyright

Document History

Date	**Version**	**Action/comments**	**Approval**
15/11/01	0.1	Document released as part of 'Initialization for New Dynamics' project	-
05/01/04	0.2	Reference to UMDP 31 added. Distinction between response functions in Fourier and normal mode space made clearer (sec. 3, para. 7). Converted to technical report and passed to Andrew Lorenc for review.	Approved for issue as a Technical Report by Andrew Lorenc
07/01/04	1.0	Issued as Forecasting Research Technical Report No. 427	-

Digital Filtering and IAU

Adam M. Clayton

7th January 2004

Abstract

In this report, we look at the theory behind three closely related techniques for removing the spurious gravity waves that result when forecasts are run from poorly balanced analyses: Digital Filtering Initialization (DFI), Digital Filtering Finalization (DFF) and Incremental Analysis Update (IAU). In particular, we show that, in a certain linear context, IAU is equivalent to the incremental form of DFF, a relationship we can exploit to improve the efficiency of incremental DFI schemes.

Contents

1 Digital Filtering Initialization (DFI) 1

2 Digital Filtering Finalization (DFF) 5

3 Incremental DFI and DFF 6

4 Incremental Analysis Update (IAU) 10

1 Digital Filtering Initialization (DFI)

DFI (Lynch and Huang, 1992) is based on the observation that the spurious gravity waves introduced by poorly balanced initial data generally have higher frequencies than are found in the balanced, meteorologically reasonable motions we are trying to predict. Thus, given a model integration contaminated by gravity wave noise, the underlying balanced evolution can be retrieved simply by filtering out frequencies above a threshold that lies between the two frequency domains. Given a time series of model states $\left\{\mathbf{x}_j = \mathbf{x}\left(t_j\right)\right\}$ valid at times

$t_j = j\Delta t$, where Δt is the model timestep, such filtering can be achieved by application of a *non-recursive digital filter* of the form

$$\mathbf{x}_n^* = \sum_{k=-M}^{+M} a_k \mathbf{x}_{n-k}, \tag{1}$$

where $\mathbf{x}_n^*$ is the time-filtered counterpart of $\mathbf{x}_n$, and the a_k are the (real) filter coefficients or weights. So that the filter coefficients are time-ordered in the same way as the model states, we may rewrite this equation in the form

$$\mathbf{x}_n^* = \sum_{k=-M}^{+M} \alpha_k \mathbf{x}_{n+k}, \tag{2}$$

where $\alpha_k = a_{-k}$. In practice, one always uses symmetrical filter coefficients — $\alpha_{-k} = \alpha_k$ — so the distinction between a_k and α_k can generally be ignored. The response of the filter is given by:

$$A(\theta) = \sum_{n=-M}^{+M} \alpha_n e^{in\theta}, \tag{3}$$

where θ is the digital frequency (radians per timestep) of a particular Fourier mode within the time-series. For symmetrical filter coefficients, $A(\theta)$ is real, implying zero phase error.

Typically, the spurious gravity waves we wish to remove have periods less than about six hours, so what is desired is a 'low-pass' filter that has a near-zero response for periods below this threshold, and a near-unit response for longer periods. Given these requirements, and a sufficiently long filter span, the theory of digital filtering can be used to choose a set of filter coefficients α_k that give something approaching the desired response. For example, Lynch and Huang, 1992 chose to use the filter shown in figure 1. An introduction to digital filtering theory, and the compromises involved in filter design, is provided in appendix B of Clayton, 2003.

The digital filter expressed in equation 2 is very simple to apply. To obtain a time-filtered state valid at timestep n, one simply integrates the model and accumulates in memory a weighted sum of model states, starting at timestep $n - M$ and ending at timestep $n + M$. In general, this time-weighting should be applied to all the prognostic variables in the model. In practice, however, some prognostic variables, such as the cloud water content, are expected to contain high frequencies even for balanced motions, and it may then be better to leave them unfiltered.

The digital filtering procedure necessarily provides time-filtered states valid *after* the start of the filter span: assuming the model integration starts at timestep zero, we have to wait until at least timestep M before we can obtain time-filtered states. Thus, if an ini-

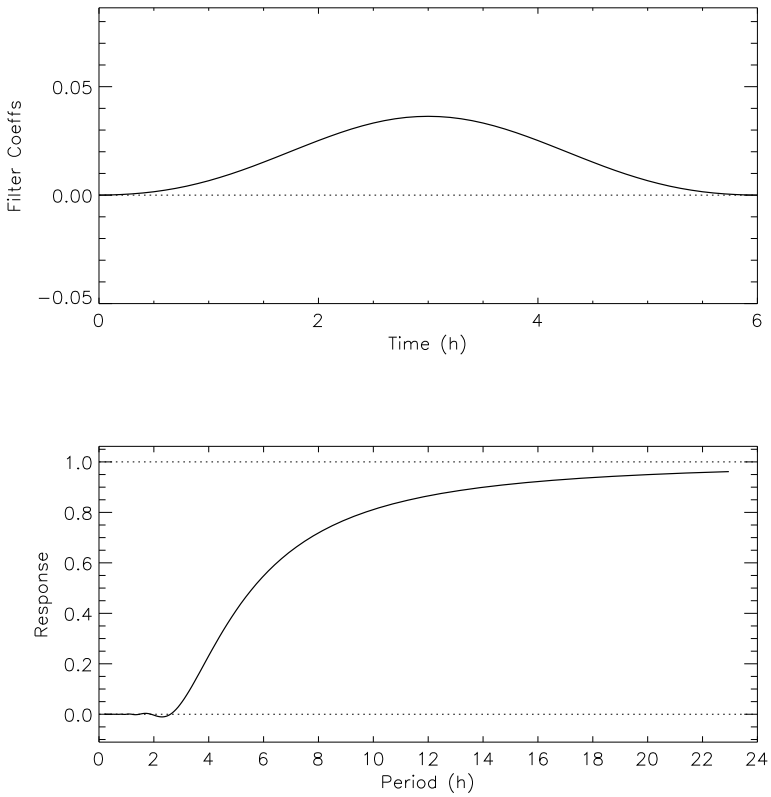

Figure 1: Coefficients and response of the filter used in Lynch and Huang, 1992: Lanczos-windowed filter, $M = 30$, $\Delta t = 6$mins, cutoff period $T_c = 6$hrs.

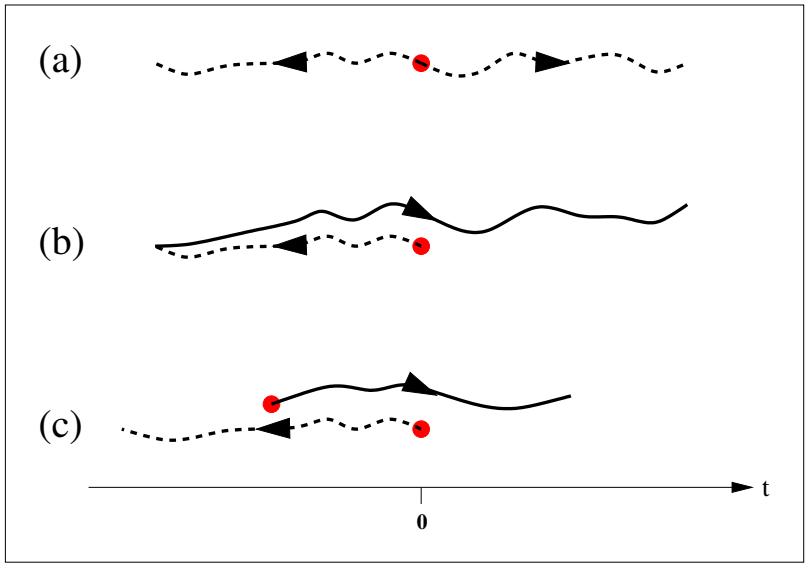

Figure 2: Schematics of three DFI schemes: (a) the adiabatic scheme of Lynch and Huang, 1992, (b) the diabatic scheme of Huang and Lynch, 1993, and (c) the diabatic scheme of Lynch et al., 1997. (Adiabatic integrations are shown as dashed lines.)

tialized analysis is required, it is necessary to combine the filtering procedure with some means of propagating states *backwards* in time to compensate for any forwards propagation. This can be achieved simply by negating the timestep of the model and disabling irreversible processes, ie. running the adiabatic component of the model backwards. In the context of a tangent linear model, Pu et al., 1997 have shown that this method provides a simple and accurate way of inverting a forwards model integration.

Given these two elements — digital filtering and backwards integrations — there are several ways one may construct a DFI scheme. The most simple is the adiabatic DFI scheme described by Lynch and Huang, 1992 and illustrated schematically in figure 2(a). Here, adiabatic integrations are run forwards and backwards from the analysis and the combined span filtered to obtain an initialized analysis. In order to incorporate diabatic effects, Huang and Lynch, 1993 suggest the scheme illustrated in figure 2(b). Here, the final state of the adiabatic backwards integration is used as the initial state for a full-physics forecast centred on the analysis time, and it is this forecast that is filtered to produce the initialized analysis. Huang and Lynch, 1993 show that this scheme has slightly better noise reduction properties than the adiabatic scheme. However, it also has some disadvantages. Firstly, because the backwards integration is adiabatic, the diabatic forwards integration will not pass back through the analysis, so the initialized analysis will inevitably include a spurious increment — the so-called diabatic discrepency. In practice, however, the discrepancy is small and is not thought to be a significant problem. Secondly, the scheme is more costly. Given that a full-physics forecast is usually about twice as costly as an adiabatic forecast,

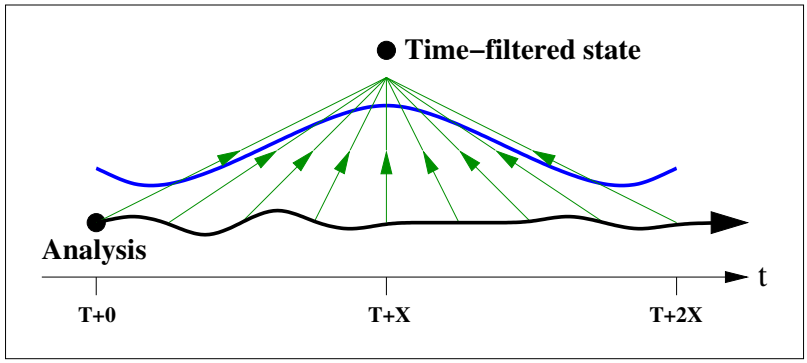

Figure 3: Schematic illustration of the DFF scheme.

the diabatic scheme costs approximately twice as much as the adiabatic scheme. Both of these problems are addressed by the scheme of Lynch et al., 1997 (figure 2(c)). Here, the idea is to filter the adiabatic backwards integration as well as the forwards integration, so that some of the gravity wave noise has been removed before the forwards integration has begun. Assuming that a similar frequency response is required in the filter's stop band, where a near-zero response is required, the length of the backwards integration generally needs to be slightly longer than before, so the scheme is a little over half as costly. Nevertheless, the scheme is just as effective, and because the forwards integration is begun from the midpoint of the backwards integration, the diabatic discrepancy is smaller.

2 Digital Filtering Finalization (DFF)

The use of adiabatic backwards integrations within DFI schemes inevitably adds spurious increments to the final initialized analysis, either through diabatic discrepancies or because backwards integrations are generally only approximate inverses of adiabatic forwards integrations. Thus, if there is no overriding requirement to have balanced fields available at the analysis time, it is usually better to exclude the backwards integrations and use a pure digital filtering scheme that gives balanced fields valid after the analysis time. This scheme, dubbed 'Digital Filtering Finalization' (DFF), has been demonstrated by Fillion et al., 1995. As they point out, within an intermittent data assimilation cycle, it is the degree of balance of the *background* fields used for generating the next analysis that is crucial; the degree of balance of any earlier states is irrelevant. As long as a balanced state can be generated in time for the production of the first set of background fields, there is nothing to be gained from having balanced fields available earlier.

The DFF procedure is illustrated schematically in figure 3. It consists simply of filter-

ing a full model forecast run directly from the analysis:

$$\mathbf{x}_M^{DFF} = \sum_{k=-M}^{+M} \alpha_k \mathbf{x}_{M+k}^a, \tag{4}$$

where $\mathbf{x}_k^a$ is the result of propagating the analysis $\mathbf{x}_0^a$ forwards k timesteps. Once the time-filtered state $\mathbf{x}_M^{DFF}$ has been obtained, the integration is restarted using $\mathbf{x}_M^{DFF}$ as initial data. The basic cost of the scheme is therefore a model integration of half the filter span, which is considerably cheaper than any comparable DFI scheme.

3 Incremental DFI and DFF

The nominal cutoff period used for DFI and DFF schemes is generally chosen to be around 6 hours. Unfortunately, as noted by Fillion et al., 1995, the limitations of filter design mean that practical filters generally have a significantly non-unit response at 12 hours, which is the period of the semi-diurnal thermal tides. For example, the filter used in Lynch and Huang, 1992, and shown in figure 1, has a response at 12 hours of 0.86, so use of this filter will lead to a slight attenuation of the semi-diurnal tides. To get around problems like this, DFI and DFF can be applied incrementally so that the background evolution to which the analysis increment is added is left unaffected. The danger in this approach is that any noise not captured by the filter will accumulate in the background forecasts. In practice, however, this does not seem to be a problem.

To apply DFI or DFF incrementally, the scheme is applied to both the analysis and the background, and then the difference taken to obtain a time-filtered increment which can be added onto the background evolution. For example, for Incremental DFF (IDFF), we have

$$\mathbf{x}_M^{IDFF} = \mathbf{x}_M^b + \sum_{k=-M}^{+M} \alpha_k \mathbf{x}_{M+k}^a - \sum_{k=-M}^{+M} \alpha_k \mathbf{x}_{M+k}^b, \tag{5}$$

or in terms of increments wrt the corresponding background states:

$$\delta\mathbf{x}_M^{IDFF} = \sum_{k=-M}^{+M} \alpha_k \delta\mathbf{x}_{M+k}^a, \tag{6}$$

where $\delta\mathbf{x}_M^{IDFF} = \mathbf{x}_M^{IDFF} - \mathbf{x}_M^b$ and $\delta\mathbf{x}_k^a = \mathbf{x}_k^a - \mathbf{x}_k^b$.

In order to develop a deeper understanding of Incremental DFI (IDFI) and IDFF, it is useful at this stage to introduce a simplified linear model of how increments behave when added onto the background evolution. With this in hand, we may then consider the fate of individual *normal modes* present in the analysis increment.

The evolution of a model state $\mathbf{x}_j$ from time t_j to time t_{j+1} is given by

$$\mathbf{x}_{j+1} = M_j\left[\mathbf{x}_j\right], \tag{7}$$

where M_j is the model operator for that timestep. The evolution of an increment $\delta\mathbf{x}_j$ defined with respect to $\mathbf{x}_j^b$ is therefore given by

$$\delta\mathbf{x}_{j+1} = M_j\left[\mathbf{x}_j^b + \delta\mathbf{x}_j\right] - M_j\left[\mathbf{x}_j^b\right]. \tag{8}$$

If $\delta\mathbf{x}_j$ is sufficiently small, we can approximate the RHS of the above equation by introducing the tangent-linear operator $\mathbf{M}_j$ given by:

$$\mathbf{M}_j = \left.\frac{\partial M_j}{\partial \mathbf{x}}\right|_{\mathbf{x}_j^b}. \tag{9}$$

We then obtain

$$\delta\mathbf{x}_{j+1} \simeq \mathbf{M}_j\delta\mathbf{x}_j. \tag{10}$$

If we now suppose that the evolution of the background state makes little difference to the evolution of the increments, ie. $\mathbf{M}_j \simeq \mathbf{M}_0$, it then follows that the evolution of increments is given approximately by

$$\delta\mathbf{x}_{j+1} = \mathbf{M}_0\delta\mathbf{x}_{j.} \tag{11}$$

Suppose now that $\mathbf{M}_0$ has eigenvectors $\mathbf{v}^l$ and corresponding (complex) eigenvalues $\lambda^l = r_l e^{i\theta_l}$, where r_l and θ_l are real. The eigenvectors of $\mathbf{M}_0$ are then the *normal modes* of equation 11. Setting $\delta\mathbf{x}_j = \mathbf{v}_j^l$, we have

$$\begin{aligned} \mathbf{v}_{j+1}^l &= \mathbf{M}_0\mathbf{v}_j^l & (12)\\ &= r_l e^{i\theta_l}\mathbf{v}_j^l, & (13) \end{aligned}$$

implying that $\mathbf{v}^l$ oscillates with digital frequency θ_l and grows by a factor of r_l every timestep.

Since the normal modes $\mathbf{v}^l$ form a basis for state space, the analysis increment $\delta\mathbf{x}_0^a$ can be written in the form

$$\delta\mathbf{x}_0^a = \sum_l c_l\mathbf{v}^l, \tag{14}$$

where the c_l are complex expansion coefficients. In the context of the linear model expressed in equation 11, the effects of schemes like IDFI and IDFF can therefore be understood in terms of their effects on individual normal modes. For example, in the context of

our linear model, equation 6 becomes

$$\delta \mathbf{x}_M^{IDFF} = \sum_{k=-M}^{+M} \alpha_k \mathbf{M}_0^{M+k} \delta \mathbf{x}_0^a. \tag{15}$$

For an individual normal mode $\mathbf{v}^l$, we then have

$$\begin{aligned} \left(\mathbf{v}_M^l\right)^{IDFF} &= \mathbf{v}_0^l \sum_{k=-M}^{+M} \alpha_k \left(r_l\right)^{M+k} e^{i(M+k)\theta_l} \\ &= \mathbf{v}_M^l \sum_{k=-M}^{+M} \alpha_k \left(r_l\right)^k e^{ik\theta_l}, \end{aligned} \tag{16}$$

so, for a given growth rate r^l, the IDFF response $R(\theta_l, r_l)$ is given by

$$R(\theta_l, r_l) = \sum_{k=-M}^{+M} \alpha_k \left(r_l\right)^k e^{ik\theta_l}. \tag{17}$$

Note that $R(\theta_l, r_l)$ expresses the response in normal mode space whereas $A(\theta)$ (equation 3) expresses the response in Fourier space. For neutral modes ($r_l = 1$), the distinction is not important as each normal mode can then be identified with a single Fourier mode. In particular, for neutral normal modes, we again see that there is no phase error as long as the filter coefficients are symmetrical. For growing or decaying normal modes, however, the correspondance between normal modes and Fourier modes is lost and the response R in normal mode space generally has an imaginary part, implying some degree of phase error. For example, for the filter shown in figure 1, the response function R for a normal mode having an e-folding time of 3 hours is shown in figure 4. We see that even for such a large growth rate, the effect on the response function, particularly its amplitude, is rather small.

Within our linear framework, IDFI and IDFF can be related to the Linear Incremental Normal Mode Initialization (NMI) scheme described by Puri et al., 1982 in which the analysis increment is projected onto a set of normal modes and the coefficients of high-frequency gravity modes set to zero, which is essentially what we are doing implicitly with IDFF and IDFI. In data assimilation experiments, Linear Incremental NMI was found to be only partially successful at removing gravity wave noise, which on the face of it might raise doubts about the schemes considered here. However, there is a crucial difference. In the scheme of Puri *et al*, the normal modes are explicitely calculated with respect to a highly simplified basic state which has no orography or any other form of geographical variation. In IDFF and IDFI, however, we are implicitly acting on normal modes defined with respect to the actual background state to which the increments are added, so the normal modes are exact. We can therefore expect these schemes to be more effective at removing noise.

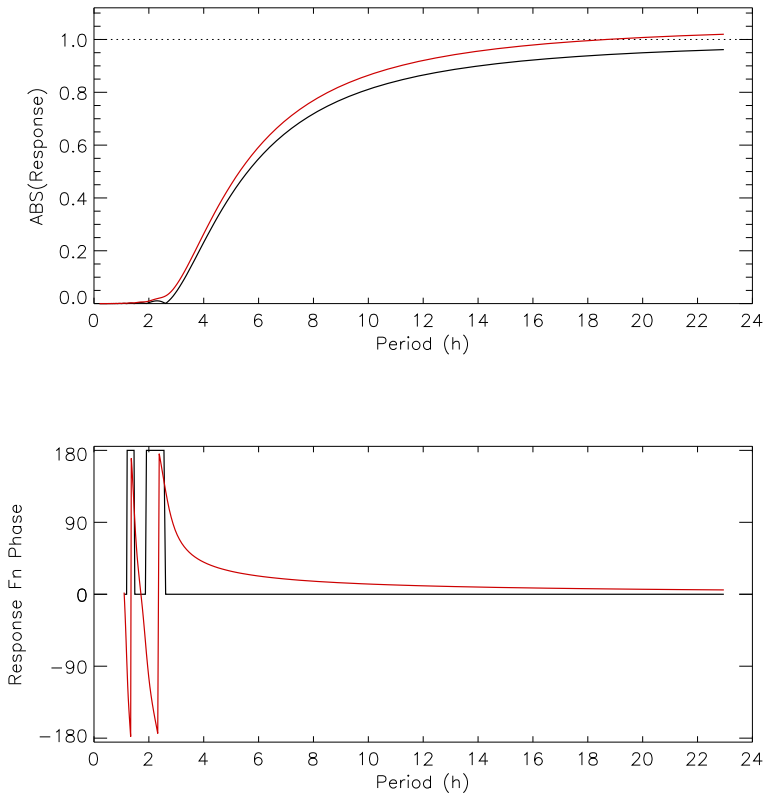

Figure 4: IDFF response, using the filter shown in figure 1, for neutral modes (black lines), and modes with an e-folding time of 3 hours (red lines).

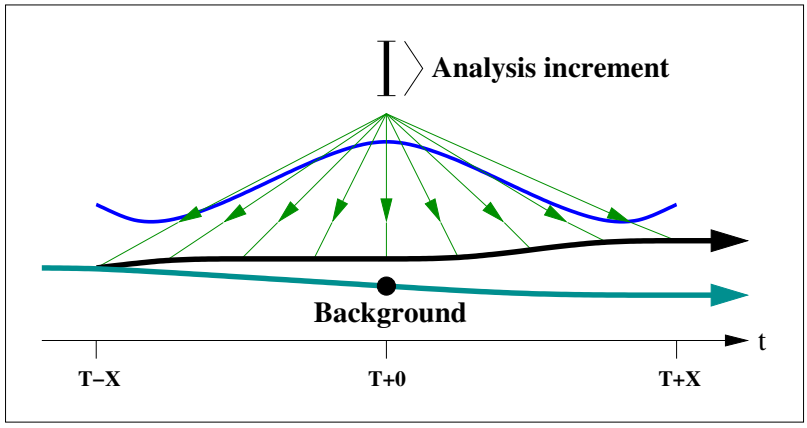

Figure 5: Schematic illustration of the IAU scheme.

4 Incremental Analysis Update (IAU)

The IAU technique (Bloom et al., 1996) consists simply of adding the analysis increment $\delta\mathbf{x}_0^a$ gradually during a short forecast centred on the analysis time rather than all in one go, as shown schematically in figure 5. At each timestep j during the insertion period, a fraction w_j of the analysis increment is added to the underlying state, so assuming the insertion period starts at timestep $-M$ and ends at timestep M we have

$$\mathbf{x}_{-M}^{IAU} = \mathbf{x}_{-M}^{b} + w_{-M}\delta\mathbf{x}_0^a, \tag{18}$$

$$\mathbf{x}_{j}^{IAU} = M_{j-1}\left[\mathbf{x}_{j-1}^{IAU}\right] + w_j\delta\mathbf{x}_0^a, \;\; -M < j \leq M, \tag{19}$$

or in terms of increments with respect to the background evolution

$$\delta\mathbf{x}_{-M}^{IAU} = w_{-M}\delta\mathbf{x}_0^a \tag{20}$$

$$\delta\mathbf{x}_{j}^{IAU} = M_{j-1}\left[\mathbf{x}_{j-1}^{b} + \mathbf{x}_{j-1}^{IAU}\right] - M_{j-1}\left[\mathbf{x}_{j-1}^{b}\right] + w_j\delta\mathbf{x}_0^a, \;\; -M < j \leq M, \tag{21}$$

The IAU scheme is best understood by assuming the IAU increments behave according to the linear model developed in the previous section (equation 11). We then have

$$\begin{aligned}
\delta\mathbf{x}^{IAU}_{-M} &= w_{-M}\delta\mathbf{x}^a_0 \\
\delta\mathbf{x}^{IAU}_{-M+1} &= \mathbf{M}_0\left(w_{-M}\delta\mathbf{x}^a\right) + w_{-M+1}\delta\mathbf{x}^a_0 \\
&= \left[w_{-M}\mathbf{M}_0 + w_{-M+1}\right]\delta\mathbf{x}^a_0 \\
\delta\mathbf{x}^{IAU}_{-M+2} &= \mathbf{M}_0\left(\left[w_{-M}\mathbf{M}_0 + w_{-M+1}\right]\delta\mathbf{x}^a_0\right) + w_{-M+2}\delta\mathbf{x}^a_0 \\
&= \left[w_{-M}\mathbf{M}_0^2 + w_{-M+1}\mathbf{M}_0 + w_{-M+2}\right]\delta\mathbf{x}^a_0 \\
&\vdots \\
\delta\mathbf{x}^{IAU}_{j} &= \sum_{k=-M}^{j} w_k\mathbf{M}_0^{j-k}\delta\mathbf{x}^a_0. \qquad (22)
\end{aligned}$$

Thus, at the end of the IAU insertion period, we have

$$\begin{aligned}
\delta\mathbf{x}^{IAU}_{M} &= \sum_{k=-M}^{+M} w_k\mathbf{M}_0^{M-k}\delta\mathbf{x}^a_0 \\
&= \sum_{k=-M}^{+M} w_{-k}\mathbf{M}_0^{M+k}\delta\mathbf{x}^a_0 \qquad (23)
\end{aligned}$$

The RHS of this equation is the same as the RHS of equation 15, except that the coefficient of $\mathbf{M}_0^{M+k}\delta\mathbf{x}^a_0$ is now w_{-k} rather than α_k. In other words, under our linear model, IDFF and IAU are equivalent except for a reversal in the time-ordering of the weights. The reason for this reversal can be understood in a straightforward manner by considering the diagram shown in figure 6. In IDFF, a mode added at time $t = 0$ and sampled at time $2X - x$, where $2X$ is the filter span, will have been evolving for a time period of $2X - x$. In the corresponding IAU scheme, in order for the mode to have been evolving for the same amount of time by the end of the insertion period, we must add the mode at time x relative to the start of the insertion period.

To summarise what we have found:

> *If the analysis increment is small, and the evolution of the background state makes little difference to the propagation of increments, an IAU scheme and an IDFF scheme having the same weights, but in reverse time-order relative to each other, are approximately equivalent.*

Because of this approximate equivalence, in determining a suitable set of weights for an IAU scheme we may use exactly the same principles as are used in designing the filter weights or coefficients for a digital filtering scheme. Since most filter design methods

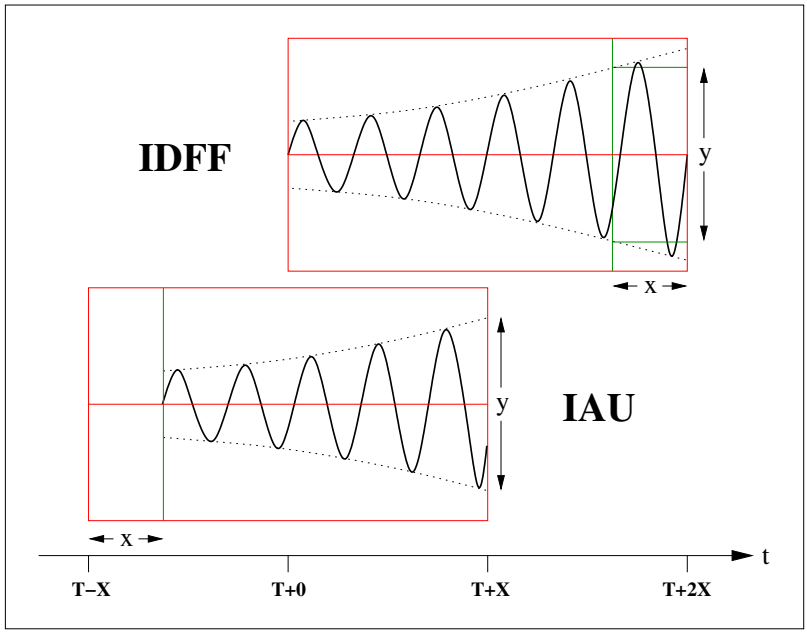

Figure 6: Illustration of why the IAU weights are reversed relative to those of the corresponding IDFF scheme.

give rise to symmetrical coefficients, the required reversal of the weights is usually a null change.

From a practical point of view, the crucial difference between IAU and IDFF is that the former is considerably less expensive computationally. For a filter span of length $2X$, the IAU scheme requires a model integration of length $2X$ to get to $T+X$, whereas IDFF requires two model integrations of the same length. In terms of integration lengths, then, IAU is half as expensive as IDFF. In addition, IDFF requires two extra model integrations to be started, which leads to further I/O and setting up costs within the model. In general, then, unless there are good reasons for doing otherwise, IAU should always be used in preference to IDFF. This recommendation applies equally to the the forwards components of incremental IDF schemes, where considerable cost savings can be made by replacing what are essentially IDFF steps with IAU integrations.

References

BLOOM, S. C., TAKACS, L. L., DA SILVA, A. M. AND LEDVINA, D. (1996). Data Assimilation Using Incremental Analysis Updates. *Mon. Wea. Rev.*, **124**, 1256–1271.

CLAYTON, A. M. (2003). Unified Model Documentation Paper 31: Incremental Analysis Update (IAU) Scheme (Model Vn. 5.5).

FILLION, L., MITCHELL, H. L., RICHIE, H. AND STANIFORTH, A. (1995). The Im-

pact of a Digital Filter Finalization Technique in a Global Data Assimilation System. *Tellus*, **47A**, 304–323.

HUANG, X.-Y. AND LYNCH, P. (1993). Diabatic Digital-Filtering Initialization: Application to the HIRLAM Model. *Mon. Wea. Rev.*, **121**, 589–603.

LYNCH, P., GIARD, D. AND IVANOVICI, V. (1997). Improving the Efficiency of a Digital Filtering Scheme for Diabatic Initialization. *Mon. Wea. Rev.*, **125**, 1976–1982.

LYNCH, P. AND HUANG, X.-Y. (1992). Initialization of the HIRLAM Model Using a Digital Filter. *Mon. Wea. Rev.*, **120**, 1019–1034.

PU, Z.-X., KALNAY, E., SELA, J. AND SZUNYOGH, I. (1997). Sensitivity of Forecast Errors to Initial Conditions with a Quasi-Inverse Linear Method. *Mon. Wea. Rev.*, **125**(10), 2479–2503.

PURI, K., BOURKE, W. AND SEAMAN, R. (1982). Incremental Linear Normal Mode Initialization in Four-Dimensional Data Assimilation. *Mon. Wea. Rev.*, **110**(12), 1773–1786.